Day : Friday
Date: 12/17/2004

PALM INTRANET

Time: 16:22:39

Inventor Name Search Result

Your Search was:

Last Name = BRINKER
First Name = BRIAN

Application#	Patent#	Status	Date Filed	Title	Inventor Name 1
09864719	Not Issued	071	05/24/2001	METHOD AND SYSTEM FOR SYSTEMATICALLY DIAGNOSING DATA PROBLEMS IN A DATABASE	BRINKER, BRIAN L.

Inventor Search Completed: No Records to Display.

Search Another: Inventor Last Name: BRINKER First Name: BRIAN Search

To go back use Back button on your browser toolbar.

Back to PALM | ASSIGNMENT | OASIS | Home page

Day : Friday
Date: 12/17/2004

PALM INTRANET

Time: 16:22:49

Inventor Name Search Result

Your Search was:

Last Name = HANSEN
First Name = BRAD

Application#	Patent#	Status	Date Filed	Title	Inventor Name 15
60170103	Not Issued	159	12/10/1999	CONDUIT AND CABLE CONSTRUCTION	HANSEN , BRADLEY S.
60169577	Not Issued	159	12/08/1999	SEDIMENT COLLECTION HEADER APPARATUS	HANSEN , BRAD
29097098	D419010	150	11/27/1998	BAG DISPENSER	HANSEN , BRAD
10944685	Not Issued	019	09/17/2004	METHOD AND APPARATUS FOR SWITCHING BETWEEN MULTIPLE WAVEFORMS	HANSEN, BRAD T.
10914554	Not Issued	020	08/09/2004	PARALLEL FILTER REALIZATION FOR WIDEBAND PROGRAMMABLE DIGITAL RADIOS	HANSEN, BRAD TERRY
10878902	Not Issued	020	06/28/2004	PARALLEL DSP DEMODULATION FOR WIDEBAND SOFTWARE-DEFINED RADIOS	HANSEN, BRAD TERRY
10422577	Not Issued	161	04/25/2003	CIRCULAR FIXED WING VTOL AIRCRAFT	HANSEN, BRAD C.
10284536	Not Issued	030	10/30/2002	ACQUISITION OF A SYNCHRONOUS CDMA TDD QPSK WAVEFORM USING POWER AND TIMING ESTIMATES	HANSEN, BRAD
10114965	Not Issued	061	04/01/2002	METHOD AND APPARATUS FOR PUNCHING PARTICLE BOARD	HANSEN, BRAD R.
09864719	Not Issued	071	05/24/2001	METHOD AND SYSTEM FOR SYSTEMATICALLY DIAGNOSING DATA PROBLEMS IN A DATABASE	HANSEN, BRAD
09740333	Not	161	12/18/2000	SEDIMENT COLLECTION	HANSEN, BRAD

	Issued			HEADER APPARATUS	
09733353	Not Issued	161	12/08/2000	CONDUIT AND CABLE CONSTRUCTION	HANSEN, BRADLEY S.
08631140	5720891	150	04/15/1996	RETRACTABLE SEDIMENT COLLECTING DEVICE FOR COVERED BASINS	HANSEN , BRAD K.
08551903	5684853	150	10/23/1995	FLEXIBLE RADIOGRAPHIC CASSETTE HOLDER	HANSEN , BRADLEY J.
08248404	Not Issued	161	05/24/1994	FLEXIBLE RADIOGRAPHIC CASSETTE HOLDER	HANSEN , BRADLEY J.

Inventor Search Completed: No Records to Display.

Search Another: Inventor **Last Name** HANSEN **First Name** BRAD Search

To go back use Back button on your browser toolbar.

Back to PALM | ASSIGNMENT | OASIS | Home page

Terms used **sql with outer join test database** Found **201** of **147,793**

Sort results by: relevance

Display results: expanded form

Save results to a Binder

Search Tips

☐ Open results in a new window

Try an Advanced Search

Try this search in The ACM Guide

Results 1 - 20 of 200 Result page: **1** 2 3 4 5 6 7 8 9 10 next

Best 200 shown Relevance scale ■■■■■

1 TEsting and reliability: Using an SQL coverage measurement for testing database applications

María José Suárez-Cabal, Javier Tuya

October 2004 **Proceedings of the 12th ACM SIGSOFT twelfth international symposium on Foundations of software engineering**

Full text available: pdf(256.74 KB) Additional Information: full citation, abstract, references, index terms

Many software applications have a component based on database management systems in which information is generally handled through SQL queries embedded in the application code. When automation of software testing is mentioned in the research, this is normally associated with programs written in imperative and structured languages. However, the problem of automated software testing applied to programs that manage databases using SQL is still an open issue. This paper presents a measurement of ...

Keywords: SQL testing, database testing, software testing, statement coverage, verification and validation

2 Query evaluation techniques for large databases

Goetz Graefe

June 1993 **ACM Computing Surveys (CSUR)**

C. J. Date
November 1984 **ACM SIGMOD Rec rd**, Volume 14 Issue 3

Full text available: pdf(2.38 MB) Additional Information: full citation, abstract, references

The ANS Database Committee (X3H2) is currently at work on a proposed standard relational database language (RDL), and has adopted as a basis for that activity a definition of the "structured query language" SQL from IBM [10]. Moreover, numerous hardware and software vendors (in addition to IBM) have already released or at least announced products that are based to a greater or lesser extent on the SQL language as defined by IBM. There can thus be little doubt that the importance of that language ...

4 Database performance in the real world: TPC-D and SAP R/3

Joachen Doppelhammer, Thomas Höppler, Alfons Kemper, Donald Kossmann
June 1997 **ACM SIGMOD Record , Proceedings of the 1997 ACM SIGMOD international conference on Management of data**, Volume 26 Issue 2

Full text available: pdf(1.54 MB) Additional Information: full citation, abstract, references, citings, index terms

Traditionally, database systems have been evaluated in isolation on the basis of standardized benchmarks (e.g., Wisconsin, TPC-C, TPC-D). We argue that very often such a performance analysis does not reflect the actual use of the DBMSs in the "real world." End users typically don't access a stand-alone database system; rather they use a comprehensive application system, in which the database system constitutes an integrated component. In order to derive performance evalu ...

5 Research sessions: spatial data: Joining interval data in relational databases

Jost Enderle, Matthias Hampel, Thomas Seidl
June 2004 **Proceedings of the 2004 ACM SIGMOD international conference on Management of data**

Full text available: pdf(552.80 KB) Additional Information: full citation, abstract, references

The increasing use of temporal and spatial data in present-day relational systems necessitates an efficient support of joins on interval-valued attributes. Standard join algorithms do not support those data types adequately, whereas special approaches for interval joins usually require an augmentation of the internal access methods which is not supported by existing relational systems. To overcome these problems we introduce new join algorithms for interval data. Based on the Relational Interval ...

6 Join processing in relational databases

Priti Mishra, Margaret H. Eich
March 1992 **ACM Computing Surveys (CSUR)**, Volume 24 Issue 1

Full text available: pdf(4.42 MB) Additional Information: full citation, abstract, references, citings, index terms, review

The join operation is one of the fundamental relational database query operations. It facilitates the retrieval of information from two different relations based on a Cartesian product of the two relations. The join is one of the most diffidult operations to implement efficiently, as no predefined links between relations are required to exist (as they are with network and hierarchical systems). The join is the only relational algebra operation that allows the combining of related tuples fro ...

Keywords: database machines, distributed processing, join, parallel processing, relational algebra

7 Optimization techniques for queries with expensive methods

Joseph M. Hellerstein

June 1998 **ACM Transacti ns n Database Systems (TODS)**, Volume 23 Issue 2

Full text available: pdf(582.16 KB) Additional Information: full citation, abstract, references, citings, ind x terms, review

Object-relational database management systems allow knowledgeable users to define new data types as well as new methods (operators) for the types. This flexibility produces an attendant complexity, which must be handled in new ways for an object-relational database management system to be efficient. In this article we study techniques for optimizing queries that contain time-consuming methods. The focus of traditional query optimizers has been on the choice of join methods and orders; selec ...

Keywords: expensive methods, extensibility, object-relational databases, predicate migration, predicate placement, query optimization

8 DB integration: Processing frequent itemset discovery queries by division and set containment join operators

Ralf Rantzau

June 2003 **Proceedings of the 8th ACM SIGMOD workshop on Research issues in data mining and knowledge discovery**

Full text available: pdf(173.83 KB) Additional Information: full citation, abstract, references, citings

SQL-based data mining algorithms are rarely used in practice today. Most performance experiments have shown that SQL-based approaches are inferior to main-memory algorithms. Nevertheless, database vendors try to integrate analysis functionalities to some extent into their query execution and optimization components in order to narrow the gap between data and processing. Such a database support is particularly important when data mining applicatons need to analyze very large datasets or when they ...

Keywords: association rule discovery, relational divison, set containment join

9 On saying "Enough already!" in SQL

Michael J. Carey, Donald Kossmann

June 1997 **ACM SIGMOD Record , Proceedings of the 1997 ACM SIGMOD international conference on Management of data**, Volume 26 Issue 2

Full text available: pdf(1.57 MB) Additional Information: full citation, abstract, references, citings, index terms

In this paper, we study a simple SQL extension that enables query writers to explicitly limit the cardinality of a query result. We examine its impact on the query optimization and run-time execution components of a relational DBMS, presenting two approaches—a Conservative approach and an Aggressive approach—to exploiting cardinality limits in relational query plans. Results obtained from an empirical study conducted using DB2 demonstrate the benefits of the SQL extensio ...

10 R* optimizer validation and performance evaluation for local queries

Lothar F. Mackert, Guy M. Lohman

June 1986 **ACM SIGMOD Record , Proceedings of the 1986 ACM SIGMOD international conference on Management of data**, Volume 15 Issue 2

Full text available: pdf(1.43 MB) Additional Information: full citation, abstract, references, citings, index terms

Few database query optimizer models have been validated against actual performance. This paper presents the methodology and results of a thorough validation of the optimizer and evaluation of the performance of the experimental distributed relational database management system R*, which inherited and extended to a distributed environment the optimization algorithms of System R. Optimizer estimated costs and actual R* resources

consumed were written ...

11 Translation with optimization from relational calculus to relational algebra having aggregate functions

Ryohei Nakano

December 1990 **ACM Transactions on Database Systems (TODS)**, Volume 15 Issue 4

Full text available: pdf(2.77 MB) Additional Information: full citation, abstract, references, citings, index terms

Most of the previous translations of relational calculus to relational algebra aimed at proving that the two languages have the equivalent expressive power, thereby generating very complicated relational algebra expressions, especially when aggregate functions are introduced. This paper presents a rule-based translation method from relational calculus expressions having both aggregate functions and null values to optimized relational algebra expressions. Thus, logical optimization is carrie ...

12 Query Language for Semantic Web: Translating XSLT programs to Efficient SQL queries

Sushant Jain, Ratul Mahajan, Dan Suciu

May 2002 **Proceedings of the eleventh international conference on World Wide Web**

Full text available: pdf(171.64 KB) Additional Information: full citation, abstract, references, citings, index terms

We present an algorithm for translating XSLT programs into SQL. Our context is that of virtual XML publishing, in which a single XML view is defined from a relational database, and subsequently queried with XSLT programs. Each XSLT program is translated into a single SQL query and run entirely in the database engine. Our translation works for a large fragment of XSLT, which we define, that includes descendant/ancestor axis, recursive templates, modes, parameters, and aggregates. We put considera ...

Keywords: SQL, XML, XSLT, query optimization, translation, virtual view

13 Research sessions: potpourri: Executing SQL over encrypted data in the database-service-provider model

Hakan Hacigümüş, Bala Iyer, Chen Li, Sharad Mehrotra

June 2002 **Proceedings of the 2002 ACM SIGMOD international conference on Management of data**

Full text available: pdf(1.25 MB) Additional Information: full citation, abstract, references, citings, index terms

Rapid advances in networking and Internet technologies have fueled the emergence of the "software as a service" model for enterprise computing. Successful examples of commercially viable software services include rent-a-spreadsheet, electronic mail services, general storage services, disaster protection services. "Database as a Service" model provides users power to create, store, modify, and retrieve data from anywhere in the world, as long as they have access to the Internet. It introduces sev ...

14 Integrating association rule mining with relational database systems: alternatives and implications

Sunita Sarawagi, Shiby Thomas, Rakesh Agrawal

June 1998 **ACM SIGMOD Record , Proceedings f the 1998 ACM SIGMOD internati nal conference n Management f data**, Volume 27 Issue 2

Full text available: pdf(2.03 MB) Additional Information: full citation, abstract, ref r nces, citings, index terms

Data mining on large data warehouses is becoming increasingly important. In support of this trend, we consider a spectrum of architectural alternatives for coupling mining with database systems. These alternatives include: loose-coupling through a SQL cursor interface; encapsulation of a mining algorithm in a stored procedure; caching the data to a file system on-the-fly and mining; tight-coupling using primarily user-defined functions; and SQL implementations for processing in the DBMS. We ...

15 Research sessions: implementation techniques: Implementing database operations using SIMD instructions

Jingren Zhou, Kenneth A. Ross

June 2002 **Proceedings of the 2002 ACM SIGMOD international conference on Management of data**

Full text available: pdf(1.39 MB) Additional Information: full citation, abstract, references, citings, index terms

Modern CPUs have instructions that allow basic operations to be performed on several data elements in parallel. These instructions are called SIMD instructions, since they apply a single instruction to multiple data elements. SIMD technology was initially built into commodity processors in order to accelerate the performance of multimedia applications. SIMD instructions provide new opportunities for database engine design and implementation. We study various kinds of operations in a database con ...

16 SchemaSQL: An extension to SQL for multidatabase interoperability

Laks V. S. Lakshmanan, Fereidoon Sadri, Subbu N. Subramanian

December 2001 **ACM Transactions on Database Systems (TODS)**, Volume 26 Issue 4

Full text available: pdf(435.89 KB) Additional Information: full citation, abstract, references, index terms, review

We provide a principled extension of SQL, called *SchemaSQL*, that offers the capability of uniform manipulation of data and schema in relational multidatabase systems. We develop a precise syntax and semantics of *SchemaSQL* in a manner that extends traditional SQL syntax and semantics, and demonstrate the following. (1) *SchemaSQL* retains the flavor of SQL while supporting querying of both data and schema. (2) It can be used to transform data in a database in a structure substa ...

Keywords: Information integration, SchemaSQL, multidatabase systems, restructuring views, schematic heterogeneity

17 XML query processing II: A comprehensive XQuery to SQL translation using dynamic interval encoding

David DeHaan, David Toman, Mariano P. Consens, M. Tamer Özsu

June 2003 **Proceedings of the 2003 ACM SIGMOD international conference on Management of data**

Full text available: pdf(242.20 KB) Additional Information: full citation, abstract, references, citings, index terms

The W3C XQuery language recommendation, based on a hierarchical and ordered document model, supports a wide variety of constructs and use cases. There is a diversity of approaches and strategies for evaluating XQuery expressions, in many cases only dealing with limited subsets of the language. In this paper we describe an implementation approach that handles XQuery with arbitrarily-nested FLWR expressions, element constructors and built-in functions (including structural comparisons). Our propos ...

18 Types and persistence in database programming languages

Malcolm P. Atkinson, O. Peter Buneman

June 1987 **ACM C mputing Surveys (CSUR)**, Volume 19 Issue 2

Full text available: pdf(7.91 MB) Additional Information: full citation, abstract, references, citings, index terms, revi w

Traditionally, the interface between a programming language and a database has either been through a set of relatively low-level subroutine calls, or it has required some form of embedding of one language in another. Recently, the necessity of integrating database and programming language techniques has received some long-overdue recognition. In response, a number of attempts have been made to construct programming languages with completely integrated database management systems. These lang ...

19 Query unnesting in object-oriented databases

Leonidas Fegaras

June 1998 **ACM SIGMOD Record , Proceedings of the 1998 ACM SIGMOD international conference on Management of data**, Volume 27 Issue 2

Full text available: pdf(1.41 MB) Additional Information: full citation, abstract, references, citings, index terms

There is already a sizable body of proposals on OODB query optimization. One of the most challenging problems in this area is query unnesting, where the embedded query can take any form, including aggregation and universal quantification. Although there is already a number of proposed techniques for query unnesting, most of these techniques are applicable to only few cases. We believe that the lack of a general and simple solution to the query unnesting problem is due to the lack of a unifo ...

20 Ad Hoc Query: a reusable database access capability

J. Wolfe

July 1994 **Proceedings of the eleventh annual Washington Ada symposium & summer ACM SIGAda meeting on Ada**

Full text available: pdf(1.06 MB) Additional Information: full citation, references, index terms

Results 1 - 20 of 200 Result page: **1** 2 3 4 5 6 7 8 9 10 next

Terms used **sql with outer join test database** Found **201** of **147,793**

Sort results by relevance Save results to a Binder

Display results expanded form Search Tips

☐ Open results in a new window

Try an Advanced Search

Try this search in The ACM Guide

Results 21 - 40 of 200 Result page: previous 1 **2** 3 4 5 6 7 8 9 10 next

Best 200 shown Relevance scale ☐■■■■

21 Database system implementation: Access path selection in a relational database management system

P. Griffiths Selinger, M. M. Astrahan, D. D. Chamberlin, R. A. Lorie, T. G. Price

May 1979 **Proceedings of the 1979 ACM SIGMOD international conference on Management of data**

Full text available: pdf(1.31 MB) Additional Information: full citation, abstract, references, citings

In a high level query and data manipulation language such as SQL, requests are stated non-procedurally, without reference to access paths. This paper describes how System R chooses access paths for both simple (single relation) and complex queries (such as joins), given a user specification of desired data as a boolean expression of predicates. System R is an experimental database management system developed to carry out research on the relational model of data. System R was designed and built b ...

22 SilkRoute: A framework for publishing relational data in XML

Mary Fernández, Yana Kadiyska, Dan Suciu, Atsuyuki Morishima, Wang-Chiew Tan

December 2002 **ACM Transactions on Database Systems (TODS)**, Volume 27 Issue 4

Full text available: pdf(687.91 KB) Additional Information: full citation, abstract, references, citings, index terms

XML is the "lingua franca" for data exchange between interenterprise applications. In this work, we describe SilkRoute, a framework for publishing relational data in XML. In SilkRoute, relational data is published in three steps: the relational tables are presented to the database administrator in a canonical XML view; the database administrator defines in the XQuery query language a public, virtual XML view over the canonical XML view; and an application formulates an XQuery query over the publ ...

Keywords: XML, XML storage systems, XQuery

23 Implementation of magic-sets in a relational database system

Inderpal Singh Mumick, Hamid Pirahesh

May 1994 **ACM SIGMOD Record , Proceedings of the 1994 ACM SIGMOD internati nal conference n Management of data**, Volume 23 Issue 2

Full text available: pdf(1.34 MB) Additional Information: full citation, abstract, referenc s, citings, index terms

We describe the implementation of the magic-sets transformation in the Starburst extensible relational database system. To our knowledge this is the first implementation of the magic-

sets transformation in a relational database system. The Starburst implementation has many novel features that make our implementation especially interesting to database practitioners (in addition to database researchers). (1) We use a cost-based heuristic for determining join orders (sips) bef ...

24 A performance evaluation of four parallel join algorithms in a shared-nothing multiprocessor environment

Donovan A. Schneider, David J. DeWitt

June 1989 **ACM SIGMOD Record , Proceedings of the 1989 ACM SIGMOD international conference on Management of data**, Volume 18 Issue 2

Full text available: pdf(1.48 MB) Additional Information: full citation, abstract, references, citings, index terms

In this paper we analyze and compare four parallel join algorithms. Grace and Hybrid hash represent the class of hash-based join methods, Simple hash represents a looping algorithm with hashing, and our last algorithm is the more traditional sort-merge. The performance of each of the algorithms with different tuple distribution policies, the addition of bit vector filters, varying amounts of main-memory for joining, and non-uniformly distributed join attribute values is studied. The Hybrid ...

25 Education/distance learning: The Kennesaw Database Courseware (KDC) with an evaluation component

Mario Guimaraes

April 2004 **Proceedings of the 42nd annual Southeast regional conference**

Full text available: pdf(582.91 KB) Additional Information: full citation, abstract, references

The Kennesaw Database Courseware (http://cofffee.kennesaw.edu) consists of animations, examples, and sample tests to support database curricula. This paper will 1) inform the recent animations/examples that have been developed, 2) discuss evaluations methods used, and 3) describe the new evaluation component that is being incorporated in the software. The work is the result of an NSF funded project (#0089412).

26 Optimizing object queries using an effective calculus

Leonidas Fegaras, David Maier

December 2000 **ACM Transactions on Database Systems (TODS)**, Volume 25 Issue 4

Full text available: pdf(641.65 KB) Additional Information: full citation, abstract, references, citings, index terms, review

Object-oriented databases (OODBs) provide powerful data abstractions and modeling facilities, but they generally lack a suitable framework for query processing and optimization. The development of an effective query optimizer is one of the key factors for OODB systems to successfully compete with relational systems, as well as to meet the performance requirements of many nontraditional applications. We propose an effective framework with a solid theoretical basis for optimizing OODB query l ...

Keywords: nested relations, object-oriented databases, query decorrelation, query optimization

27 Parallelism and its price: a case study of nonstop SQL/MP

Susanne Englert, Ray Glasstone, Waqar Hasan

December 1995 **ACM SIGMOD Record**, Volume 24 Issue 4

Full text available: pdf(1.09 MB) Additional Information: full citation, abstract, index terms

We describe the use of parallel execution techniques and measure the price of parallel execution in NonStop SQL/MP, a commercial parallel database system from Tandem

Computers. Non-Stop SQL uses intra-operator parallelism to parallelize joins, groupings and scans. Parallel execution consists of starting up several processes and communicating data between them. Our measurements show (a) Startup costs are negligible when processes are reused rather than created afresh (b) Communication costs ...

28 ANSI SQL hierarchical processing can fully integrate native XML

Michael M. David
March 2003 **ACM SIGMOD Record**, Volume 32 Issue 1

Full text available: pdf(162.35 KB) Additional Information: full citation, abstract, references

Most SQL-based XML vendor support is through interoperation and not integration. One reason for this is that XML is inherently hierarchical and SQL is supposedly not. This paper demonstrates how ANSI SQL along with its relational Cartesian product model can naturally perform complete and flexible hierarchical query processing. With this ANSI SQL inherent hierarchical processing capability, native XML data can be fully and seamlessly integrated into SQL processing and operated on at a full hierar ...

29 Optimization of nested SQL queries revisited

Richard A. Ganski, Harry K. T. Wong
December 1987 **ACM SIGMOD Record , Proceedings of the 1987 ACM SIGMOD international conference on Management of data**, Volume 16 Issue 3

Full text available: pdf(1.08 MB) Additional Information: full citation, abstract, references, citings, index terms

Current methods of evaluating nested queries in the SQL language can be inefficient in a variety of query and data base contexts. Previous research in the area of nested query optimization which sought methods of reducing evaluation costs is summarized, including a classification scheme for nested queries, algorithms designed to transform each type of query to a logically equivalent form which may then be evaluated more efficiently, and a description of a major bug in one of these algorithm ...

30 Efficient evaluation of XML middle-ware queries

Mary Fernandez, Atsuyuki Morishima, Dan Suciu
May 2001 **ACM SIGMOD Record , Proceedings of the 2001 ACM SIGMOD international conference on Management of data**, Volume 30 Issue 2

Full text available: pdf(414.15 KB) Additional Information: full citation, abstract, references, citings, index terms

We address the problem of efficiently constructing materialized XML views of relational databases. In our setting, the XML view is specified by a query in the declarative query language of a middle-ware system, called SilkRoute. The middle-ware system evaluates a query by sending one or more SQL queries to the target relational database, integrating the resulting tuple streams, and adding the XML tags. We focus on how to best choose the SQL queries, without having control over the target RDBM ...

31 The BUCKY object-relational benchmark

Michael J. Carey, David J. DeWitt, Jeffrey F. Naughton, Mohammad Asgarian, Paul Brown, Johannes E. Gehrke, Dhaval N. Shah
June 1997 **ACM SIGMOD Record , Proceedings of the 1997 ACM SIGMOD international conference on Management of data**, Volume 26 Issue 2

Full text available: pdf(1.48 MB) Additional Information: full citation, abstract, references, citings, index terms

According to various trade journals and corporate marketing machines, we are now on the verge of a revolution—the object-relational database revolution. Since we believe that no one should face a revolution without appropriate armaments, this paper presents BUCKY, a new benchmark for object-relational database systems. BUCKY is a query-oriented

benchmark that tests many of the key features offered by object-relational systems, including row types and inheritance, references and path e ...

32 Video analysis, retrieval, and summarizing: Video query processing in the VDBMS testbed for video database research

Walid Aref, Moustafa Hammad, Ann Christine Catlin, Ihab Ilyas, Thanaa Ghanem, Ahmed Elmagarmid, Mirette Marzouk

November 2003 **Proceedings of the 1st ACM international workshop on Multimedia databases**

Full text available: pdf(357.93 KB) Additional Information: full citation, abstract, references, index terms

The increased use of video data sets for multimedia-based applications has created a demand for strong video database support, including efficient methods for handling the content-based query and retrieval of video data. Video query processing presents significant research challenges, mainly associated with the size, complexity and unstructured nature of video data. A video query processor must support video operations for search by content and streaming, new query types, and the incorporation o ...

Keywords: continuous query, query processing, rank-join algorithm, stream processing, video database, window-join algorithm

33 Query Optimization in Database Systems

Matthias Jarke, Jurgen Koch

June 1984 **ACM Computing Surveys (CSUR)**, Volume 16 Issue 2

Full text available: pdf(2.84 MB) Additional Information: full citation, references, citings, index terms

34 Research sessions: XML I: Storing and querying ordered XML using a relational database system

Igor Tatarinov, Stratis D. Viglas, Kevin Beyer, Jayavel Shanmugasundaram, Eugene Shekita, Chun Zhang

June 2002 **Proceedings of the 2002 ACM SIGMOD international conference on Management of data**

Full text available: pdf(1.38 MB) Additional Information: full citation, abstract, references, citings, index terms

XML is quickly becoming the *de facto* standard for data exchange over the Internet. This is creating a new set of data management requirements involving XML, such as the need to store and query XML documents. Researchers have proposed using relational database systems to satisfy these requirements by devising ways to "shred" XML documents into relations, and translate XML queries into SQL queries over these relations. However, a key issue with such an approach, which has largely been ignor ...

35 Extensible query processing in starburst

L. M. Haas, J. C. Freytag, G. M. Lohman, H. Pirahesh

June 1989 **ACM SIGMOD Record , Proceedings of the 1989 ACM SIGMOD international conference on Management of data**, Volume 18 Issue 2

Full text available: pdf(1.63 MB) Additional Information: full citation, abstract, references, citings, index terms

Today's DBMSs are unable to support the increasing demands of the various applications that would like to use a DBMS. Each kind of application poses new requirements for the DBMS. The Starburst project at IBM's Almaden Research Center aims to extend relational DBMS technology to bridge this gap between applications and the DBMS. While providing a full function relational system to enable sharing across applications, Starburst will also allow

(sophisticated) programmers to add many kinds of ...

36 Session 4: Processing queries with quantifiers a horticultural approach
Umeshwar Dayal
March 1983 **Pr ceedings f the 2nd ACM SIGACT-SIGMOD symposium n Principles of database systems**
Full text available: pdf(1.11 MB) Additional Information: full citation, abstract, references, citings

Most research on query processing has focussed on quantifier-free conjunctive queries. Existing techniques for processing queries with quantifiers either compile the query into a nested loop program or use variants of Codd's reduction from the Relational Calculus to the Relational Algebra. In this paper we propose an alternative technique that uses an algebra of graft and prune operations on trees. This technique provides a significant savings in space and time. We show how to transform a quanti ...

37 Full papers: Horizontal aggregations for building tabular data sets
Carlos Ordonez
June 2004 **Proceedings of the 9th ACM SIGMOD workshop on Research issues in data mining and knowledge discovery**
Full text available: pdf(183.18 KB) Additional Information: full citation, abstract, references

In a data mining project, a significant portion of time is devoted to building a data set suitable for analysis. In a relational database environment, building such data set usually requires joining tables and aggregating columns with SQL queries. Existing SQL aggregations are limited since they return a single number per aggregated group, producing one row for each computed number. These aggregations help, but a significant effort is still required to build data sets suitable for data mining pu ...

38 Research sessions: query optimization: Canonical abstraction for outerjoin optimization
Jun Rao, Hamid Pirahesh, Calisto Zuzarte
June 2004 **Proceedings of the 2004 ACM SIGMOD international conference on Management of data**
Full text available: pdf(197.25 KB) Additional Information: full citation, abstract, references

Outerjoins are an important class of joins and are widely used in various kinds of applications. It is challenging to optimize queries that contain outerjoins because outerjoins do not always commute with inner joins. Previous work has studied this problem and provided techniques that allow certain reordering of the join sequences. However, the optimization of outerjoin queries is still not as powerful as that of inner joins.An inner join query can always be canonically represented as a sequence ...

39 Fast algorithms for universal quantification in large databases
Goetz Graefe, Richard L. Cole
June 1995 **ACM Transactions on Database Systems (TODS)**, Volume 20 Issue 2
Full text available: pdf(3.51 MB) Additional Information: full citation, abstract, references, citings, index terms, review

Universal quantification is not supported directly in most database systems despite the fact that it adds significant power to a system's query processing and inference capabilities, in particular for the analysis of many-to-many relationships and of set-valued attributes. One of the main reasons for this omission has been that universal quantification algorithms and their performance have not been explored for large databases. In this article, we describe and compare three known algorithms ...

40 A high-performance parallel database architecture
C. H. C. Leung, H. T. Ghogomu

August 1993 **Pr ceedings f the 7th internati nal c nference on Superc mputing**

Full text available: pdf(911.37 KB) Additional Information: full citation, abstract, references, citings, ind x terms

A high-performance parallel system for processing databases is presented, which adopts a distributed memory architecture and has been successfully implemented on a transputer platform. In addition to developing and implementing a variety of rules and schemes for parallelizing database queries, general analytic models for distributed memory database processing have been formulated, which have been successfully validated against measurements. Experimental data also indicate that the system is ...

Results 21 - 40 of 200 Result page: previous 1 **2** 3 4 5 6 7 8 9 10 next

Terms used **sql with outer join test database** Found **201** of **147,793**

Sort results by: relevance
Display results: expanded form
Save results to a Binder
Search Tips
☐ Open results in a new window

Try an Advanced Search
Try this search in The ACM Guide

Results 41 - 60 of 200 Result page: previous 1 2 **3** 4 5 6 7 8 9 10 next
Best 200 shown Relevance scale

41 XML query processing II: On relational support for XML publishing: beyond sorting and tagging
Surajit Chaudhuri, Raghav Kaushik, Jeffrey F. Naughton
June 2003 **Proceedings of the 2003 ACM SIGMOD international conference on Management of data**
Full text available: pdf(237.72 KB) Additional Information: full citation, abstract, references, index terms

In this paper, we study whether the need for efficient XML publishing brings any new requirements for relational query engines, or if sorting query results in the relational engine and tagging them in middleware is sufficient. We observe that the mismatch between the XML data model and the relational model requires relational engines to be enhanced for efficiency. Specifically, they need to support relation valued variables. We discuss how such support can be provided through the addition of an ...

42 DB-5 (databases): potpourri: Computing consistent query answers using conflict hypergraphs
Jan Chomicki, Jerzy Marcinkowski, Slawomir Staworko
November 2004 **Proceedings of the Thirteenth ACM conference on Information and knowledge management**
Full text available: pdf(261.96 KB) Additional Information: full citation, abstract, references, index terms

A consistent query answer in a possibly inconsistent database is an answer which is true in every (minimal) repair of the database. We present here a practical framework for computing consistent query answers for large, possibly inconsistent relational databases. We consider relational algebra queries without projection, and denial constraints. Because our framework handles union queries, we can effectively (and efficiently) extract indefinite disjunctive information from an inconsistent data ...

Keywords: inconsistency, integrity constraints, query processing

43 Multidatabase systems: Exploiting uniqueness in query optimization
G. N. Paulley, Per-Åke Larson
October 1993 **Proceedings of the 1993 conference f the Centre f r Advanced Studies on C llaborative research: distributed c mputing - V lume 2**
Full text available: pdf(1.27 MB) Additional Information: full citation, abstract, r ferences

Consider an SQL query that specifies duplicate elimination via a DISTINCT clause. Because

duplicate elimination often requires an expensive sort of the query result, it is often worthwhile to identify situations where the DISTINCT clause is unnecessary, to avoid the sort altogether. We prove a necessary and sufficient condition for deciding if a query requires duplicate elimination. The condition exploits knowledge about keys, table constraints, and query predicates. Because the condition cannot ...

44 Research sessions: consistency and availability: Relaxed currency and consistency: how to say "good enough" in SQL

Hongfei Guo, Per-Åke Larson, Raghu Ramakrishnan, Jonathan Goldstein

June 2004 **Proceedings of the 2004 ACM SIGMOD international conference on Management of data**

Full text available: pdf(606.81 KB) Additional Information: full citation, abstract, references, citings

Despite the widespread and growing use of asynchronous copies to improve scalability, performance and availability, this practice still lacks a firm semantic foundation. Applications are written with some understanding of which queries can use data that is not entirely current and which copies are "good enough"; however, there are neither explicit requirements nor guarantees. We propose to make this knowledge available to the DBMS through explicit currency and consistency (C&C) constraints in qu ...

45 Towards an efficient evaluation of general queries: quantifier and disjunction processing revisited

Francois Bry

June 1989 **ACM SIGMOD Record , Proceedings of the 1989 ACM SIGMOD international conference on Management of data**, Volume 18 Issue 2

Full text available: pdf(1.63 MB) Additional Information: full citation, abstract, references, citings, index terms

Database applications often require to evaluate queries containing quantifiers or disjunctions, e.g., for handling general integrity constraints. Existing efficient methods for processing quantifiers depart from the relational model as they rely on non-algebraic procedures. Looking at quantified query evaluation from a new angle, we propose an approach to process quantifiers that makes use of relational algebra operators only. Our approach performs in two phases. The first phase nor ...

46 NSF workshop on industrial/academic cooperation in database systems

Mike Carey, Len Seligman

March 1999 **ACM SIGMOD Record**, Volume 28 Issue 1

Full text available: pdf(1.96 MB) Additional Information: full citation, index terms

47 Extending a database system with procedures

Michael Stonebraker, Jeff Anton, Eric Hanson

September 1987 **ACM Transactions on Database Systems (TODS)**, Volume 12 Issue 3

Full text available: pdf(2.15 MB) Additional Information: full citation, abstract, references, citings, index terms, review

This paper suggests that more powerful database systems (DBMS) can be built by supporting database procedures as full-fledged database objects. In particular, allowing fields of a database to be a collection of queries in the query language of the system is shown to allow the natural expression of complex data relationships. Moreover, many of the features present in object-oriented systems and semantic data models can be supported by this facility. In order to implement this cons ...

48 Research sessions: query progress: Toward a progress indicator for database queries

Gang Luo, Jeffrey F. Naughton, Curt J. Ellmann, Michael W. Watzke
June 2004 **Pr ceedings of the 2004 ACM SIGMOD international c nference n Management of data**
Full text available: pdf(228.58 KB) Additional Information: full citation, abstract, references

Many modern software systems provide progress indicators for long-running tasks. These progress indicators make systems more user-friendly by helping the user quickly estimate how much of the task has been completed and when the task will finish. However, none of the existing commercial RDBMSs provides a non-trival progress indicator for long-running queries. In this paper, we consider the problem of supporting such progress indicators. After discussing the goals and challenges inherent in this ...

49 Research sessions: non-standard query processing: Optimization of query streams using semantic prefetching
Ivan T. Bowman, Kenneth Salem
June 2004 **Proceedings of the 2004 ACM SIGMOD international conference on Management of data**
Full text available: pdf(224.40 KB) Additional Information: full citation, abstract, references

Streams of relational queries submitted by client applications to database servers contain patterns that can be used to predict future requests. We present the Scalpel system, which detects these patterns and optimizes request streams using context-based predictions of future requests. Scalpel uses its predictions to provide a form of semantic prefetching, which involves combining a predicted series of requests into a single request that can be issued immediately. Scalpel's semantic prefetching ...

50 An algebraic approach to static analysis of active database rules
Elena Baralis, Jennifer Widom
September 2000 **ACM Transactions on Database Systems (TODS)**, Volume 25 Issue 3
Full text available: pdf(391.93 KB) Additional Information: full citation, abstract, references, citings, index terms

Rules in active database systems can be very difficult to program due to the unstructured and unpredictable nature of rule processing. We provide static analysis techniques for predicting whether a given rule set is guaranteed to terminate and whether rule execution is confluent (guaranteed to have a unique final state). Our methods are based on previous techniques for analyzing rules in active database systems. We improve considerably on the previous techniques by providing analysis criter ...

Keywords: active database systems, confluence, database rule processing, database trigger processing, termination

51 Optimizing queries using materialized views: a practical, scalable solution
Jonathan Goldstein, Per-Åke Larson
May 2001 **ACM SIGMOD Record , Proceedings of the 2001 ACM SIGMOD international conference on Management of data**, Volume 30 Issue 2
Full text available: pdf(202.08 KB) Additional Information: full citation, abstract, references, citings, index terms, review

Materialized views can provide massive improvements in query processing time, especially for aggregation queries over large tables. To realize this potential, the query optimizer must know how and when to exploit materialized views. This paper presents a fast and scalable algorithm for determining whether part or all of a query can be computed from materialized views and describes how it can be incorporated in transformation-based optimizers. The current version handles views composed of sele ...

Keywords: materialized views, query optimization, view matching

52 Automatic high-quality reengineering of database programs by abstraction, transformation and reimplementation

Yossi Cohen, Yishai A. Feldman

July 2003 **ACM Transactions on Software Engineering and Methodology (TOSEM)**, Volume 12 Issue 3

Full text available: pdf(245.97 KB) Additional Information: full citation, abstract, references, index terms

Old-generation database models, such as the indexed-sequential, hierarchical, or network models, provide record-level access to their data, with all application logic residing in the hosting program. In contrast, relational databases can perform complex operations, such as filter, aggregation, and join, on multiple records without an external specification of the record-access logic. Programs written for relational databases attempt to move as much of the application logic as possible into the d ...

Keywords: Database program reengineering, query graphs, temporal abstraction, the plan calculus

53 Industrial track session 6: CRM and query optimization: WinMagic: subquery elimination using window aggregation

Calisto Zuzarte, Hamid Pirahesh, Wenbin Ma, Qi Cheng, Linqi Liu, Kwai Wong

June 2003 **Proceedings of the 2003 ACM SIGMOD international conference on Management of data**

Full text available: pdf(187.57 KB) Additional Information: full citation, abstract, references, index terms

Database queries often take the form of correlated SQL queries. Correlation refers to the use of values from the outer query block to compute the inner subquery. This is a convenient paradigm for SQL programmers and closely mimics a function invocation paradigm in a typical computer programming language. Queries with correlated subqueries are also often created by SQL generators that translate queries from application domain-specific languages into SQL. Another significant class of queries that ...

54 A declarative approach to optimize bulk loading into databases

Sihem Amer-Yahia, Sophie Cluet

June 2004 **ACM Transactions on Database Systems (TODS)**, Volume 29 Issue 2

Full text available: pdf(1.00 MB) Additional Information: full citation, abstract, references, index terms

Applications, such as warehouse maintenance, need to load large data volumes regularly. The efficiency of loading depends on the resources that are available at the source and at the target systems. Our work aims to understand the performance criteria that are involved in bulk loading data into a database and to devise tailored optimization strategies.Unlike commercial systems and previous research on the same topic, our approach follows the fundamental database principle of physical-logical ind ...

Keywords: Declarative bulk loading, algebra, recovery, side-effects

55 A cognitive model of database querying: a tool for novice instruction

M. S. Schlager, W. C. Ogden

April 1986 **ACM SIGCHI Bulletin , Proceedings of the SIGCHI conference n Human factors in computing systems**, Volume 17 Issue 4

Full text available: pdf(804.99 KB) Additional Information: full citation, abstract, references, index terms

Two experiments examine the effects of incorporating user knowledge into the design of

training materials for a database querying system. In Experiment 1 an informal cognitive model of a query language is derived from the verbal reports of expert users, and incorporated into existing documentation. Two groups of subjects were then asked to solve queries using either the revised or original manual. In Experiment 2 the cognitive model was formalized to explicitly describe the conceptual and p ...

56 Special issue on prototypes of deductive database systems: The CORAL deductive system

Raghu Ramakrishnan, Divesh Srivastava, S. Sudarshan, Praveen Seshadri

April 1994 **The VLDB Journal — The International Journal on Very Large Data Bases**, Volume 3 Issue 2

Full text available: pdf(3.03 MB) Additional Information: full citation, abstract, references, citings

CORAL is a deductive system that supports a rich declarative language, and an interface to C++, which allows for a combination of declarative and imperative programming. A CORAL declarative program can be organized as a collection of interacting modules. CORAL supports a wide range of evaluation strategies, and automatically chooses an efficient strategy for each module in the program. Users can guide query optimization by selecting from a wide range of control choices. The CORAL system provides ...

Keywords: deductive database, logic programming system, query language

57 Outerjoin simplification and reordering for query optimization

César Galindo-Legaria, Arnon Rosenthal

March 1997 **ACM Transactions on Database Systems (TODS)**, Volume 22 Issue 1

Full text available: pdf(616.62 KB) Additional Information: full citation, references, citings, index terms, review

Keywords: outerjoins, query optimization, query reordering

58 Reusing invariants: a new strategy for correlated queries

Jun Rao, Kenneth A. Ross

June 1998 **ACM SIGMOD Record , Proceedings of the 1998 ACM SIGMOD international conference on Management of data**, Volume 27 Issue 2

Full text available: pdf(1.55 MB) Additional Information: full citation, abstract, references, citings, index terms

Correlated queries are very common and important in decision support systems. Traditional nested iteration evaluation methods for such queries can be very time consuming. When they apply, query rewriting techniques have been shown to be much more efficient. But query rewriting is not always possible. When query rewriting does not apply, can we do something better than the traditional nested iteration methods? In this paper, we propose a new invariant technique to evaluate correlated queries ...

59 Special issue on spatial database systems: An introduction to spatial database systems

Ralf Hartmut Güting

October 1994 **The VLDB Journal — The International Journal on Very Large Data Bases**, Volume 3 Issue 4

Full text available: pdf(2.50 MB) Additional Information: full citation, abstract, references, citings

We propose a definition of a spatial database system as a database system that offers spatial data types in its data model and query language, and supports spatial data types in its implementation, providing at least spatial indexing and spatial join methods. Spatial

database systems offer the underlying database technology for geographic information systems and other applications. We survey data modeling, querying, data structures and algorithms, and system architecture for such systems. The em ...

60 Industry session 3: data analysis, mining, and managing XML: SQL text parsing for information retrieval

David Holmes

November 2003 **Proceedings of the twelfth international conference on Information and knowledge management**

Full text available: pdf(175.06 KB) Additional Information: full citation, abstract, references, index terms

The concept of using a relational database to perform information retrieval (IR) search functions is well established. Prior work demonstrates the capability to perform common functions and advanced ranking algorithms using standard, unchanged SQL. The previous work does not address the preprocessing of unstructured text within the relational model. In fact, the parsing of the unstructured data into a structured data set was done outside of the database, usually using sequential programming lang ...

Keywords: SQL, text parsing

Results 41 - 60 of 200 Result page: previous 1 2 **3** 4 5 6 7 8 9 10 next

Terms used **sql with outer join test database**

Found **201** of **147,793**

Sort results by: relevance

Display results: expanded form

Save results to a Binder

Search Tips

☐ Open results in a new window

Try an Advanced Search

Try this search in The ACM Guide

Results 61 - 80 of 200 Result page: previous 1 2 3 **4** 5 6 7 8 9 10 next

Best 200 shown

Relevance scale

61 Predicate migration: optimizing queries with expensive predicates

Joseph M. Hellerstein, Michael Stonebraker

June 1993 **ACM SIGMOD Record , Proceedings of the 1993 ACM SIGMOD international conference on Management of data**, Volume 22 Issue 2

independence, parameter markers, and so on. Cardinality estimation errors may cause the optimizer to choose a sub-optimal plan. We present an approach to query processing that is extremely robust because it is able to detect and re ...

64 The state of the art in distributed query processing

Donald Kossmann

December 2000 **ACM Computing Surveys (CSUR)**, Volume 32 Issue 4

Full text available: pdf(455.39 KB) Additional Information: full citation, abstract, references, citings, index terms

Distributed data processing is becoming a reality. Businesses want to do it for many reasons, and they often must do it in order to stay competitive. While much of the infrastructure for distributed data processing is already there (e.g., modern network technology), a number of issues make distributed data processing still a complex undertaking: (1) distributed systems can become very large, involving thousands of heterogeneous sites including PCs and mainframe server machines; (2) the stat ...

Keywords: caching, client-server databases, database application systems, dissemination-based information systems, economic models for query processing, middleware, multitier architectures, query execution, query optimization, replication, wrappers

65 A transformation-based approach to optimizing loops in database programming languages

Daniel F. Lieuwen, David J. DeWitt

June 1992 **ACM SIGMOD Record , Proceedings of the 1992 ACM SIGMOD international conference on Management of data**, Volume 21 Issue 2

Full text available: pdf(1.21 MB) Additional Information: full citation, abstract, references, citings, index terms

Database programming languages like O2, E, and O++ include the ability to iterate through a set. Nested iterators can be used to express joins. This paper describes compile-time optimizations similar to relational transformations like join reordering for such programming constructs. This paper also shows how to use a standard transformation-based optimizer to optimize these joins. An optimizer built using the EXODUS Opt ...

66 Ripple joins for online aggregation

Peter J. Haas, Joseph M. Hellerstein

June 1999 **ACM SIGMOD Record , Proceedings of the 1999 ACM SIGMOD international conference on Management of data**, Volume 28 Issue 2

Full text available: pdf(1.78 MB) Additional Information: full citation, abstract, references, citings, index terms

We present a new family of join algorithms, called ripple joins, for online processing of multi-table aggregation queries in a relational database management system (DBMS). Such queries arise naturally in interactive exploratory decision-support applications. Traditional offline join algorithms are designed to minimize the time to completion of the query. In contrast, ripple joins are designed to minimize the time until an acceptably precise estimate of the query result is availa ...

67 On the use of distributed joins for processing interlibrary loans

J. Michael Bennett, P. Neo

April 1992 **Proceedings of the 1992 ACM/SIGAPP Symp sium on Applied c mputing: technol gical challenges of the 1990's**

Full text available: pdf(1.09 MB) Additional Information: full citation, references, index terms

68 Office-by-example: an integrated office system and database manager
Kyu-Young Whang, Art Ammann, Anthony Bolmarcich, Maria Hanrahan, Guy Hochgesang, Kuan-Tsae Huang, Al Khorasani, Ravi Krishnamurthy, Gary Sockut, Paula Sweeney, Vance Waddle, Moshé Zloof
October 1987 **ACM Transactions on Information Systems (TOIS)**, Volume 5 Issue 4

Full text available: pdf(2.86 MB) Additional Information: full citation, abstract, references, citings, index terms, review

Office-by-Example (OBE) is an integrated office information system that has been under development at IBM Research. OBE, an extension of Query-by-Example, supports various office features such as database tables, word processing, electronic mail, graphics, images, and so forth. These seemingly heterogeneous features are integrated through a language feature called example elements. Applications involving example elements are processed by the database manager, an integrated ...

69 Fundamental techniques for order optimization
David Simmen, Eugene Shekita, Timothy Malkemus
June 1996 **ACM SIGMOD Record , Proceedings of the 1996 ACM SIGMOD international conference on Management of data**, Volume 25 Issue 2

Full text available: pdf(1.07 MB) Additional Information: full citation, abstract, references, citings, index terms

Decision support applications are growing in popularity as more business data is kept on-line. Such applications typically include complex SQL queries that can test a query optimizer's ability to produce an efficient access plan. Many access plan strategies exploit the physical ordering of data provided by indexes or sorting. Sorting is an expensive operation, however. Therefore, it is imperative that sorting is optimized in some way or avoided all together. Toward that goal, this paper describe ...

70 Practical predicate placement
Joseph M. Hellerstein
May 1994 **ACM SIGMOD Record , Proceedings of the 1994 ACM SIGMOD international conference on Management of data**, Volume 23 Issue 2

Full text available: pdf(1.22 MB) Additional Information: full citation, abstract, references, citings, index terms

Recent work in query optimization has addressed the issue of placing expensive predicates in a query plan. In this paper we explore the predicate placement options considered in the Montage DBMS, presenting a family of algorithms that form successively more complex and effective optimization solutions. Through analysis and performance measurements of Montage SQL queries, we classify queries and highlight the simplest solution that will optimize each class correctly. We demonstrate limitatio ...

71 Online analytic processing (OLAP): Spreadsheets in RDBMS for OLAP
Andrew Witkowski, Srikanth Bellamkonda, Tolga Bozkaya, Gregory Dorman, Nathan Folkert, Abhinav Gupta, Lei Shen, Sankar Subramanian
June 2003 **Proceedings of the 2003 ACM SIGMOD international conference on Management of data**

Full text available: pdf(182.22 KB) Additional Information: full citation, abstract, references, index terms

One of the critical deficiencies of SQL is lack of support for n-dimensional array-based computations which are frequent in OLAP environments. Relational OLAP (ROLAP) applications have to emulate them using joins, recently introduced SQL Window Functions [18] and complex and inefficient CASE expressions. The designated place in SQL for specifying calculations is the SELECT clause, which is extremely limiting and forces the user to generate queries using nested views, subqueries and complex joins ...

72 Statistical profile estimation in database systems
Michael V. Mannino, Paicheng Chu, Thomas Sager
September 1988 **ACM Computing Surveys (CSUR)**, Volume 20 Issue 3

Full text available: pdf(2.94 MB) Additional Information: full citation, abstract, references, citings, index terms

A statistical profile summarizes the instances of a database. It describes aspects such as the number of tuples, the number of values, the distribution of values, the correlation between value sets, and the distribution of tuples among secondary storage units. Estimation of database profiles is critical in the problems of query optimization, physical database design, and database performance prediction. This paper describes a model of a database of profile, relates this model to estimating ...

73 Relational database: a practical foundation for productivity
E. F. Codd
February 1982 **Communications of the ACM**, Volume 25 Issue 2

Full text available: pdf(899.90 KB) Additional Information: full citation, references, citings, index terms

Keywords: data integrity, data manipulation, data structure, database, productivity, relational database, relational model

74 Heraclitus: elevating deltas to be first-class citizens in a database programming language
Shahram Ghandeharizadeh, Richard Hull, Dean Jacobs
September 1996 **ACM Transactions on Database Systems (TODS)**, Volume 21 Issue 3

Full text available: pdf(3.76 MB) Additional Information: full citation, abstract, references, citings, index terms, review

Traditional database systems provide a user with the ability to query and manipulate one database state, namely the current database state. However, in several emerging applications, the ability to analyze "what-if" scenarios in order to reason about the impact of an update (before committing that update) is of paramount importance. Example applications include hypothetical database access, active database management systems, and version management, to name a few. The central th ...

Keywords: active databases, deltas, execution model for rule application, hypothetical access, hypothetical database state

75 Array-driven simulation of real databases
William S. Keezer
December 1998 **Proceedings of the 30th conference on Winter simulation**

Full text available: pdf(70.92 KB) Additional Information: full citation, references, index terms

76 Query processing: Estimating compilation time of a query optimizer
Ihab F. Ilyas, Jun Rao, Guy Lohman, Dengfeng Gao, Eileen Lin
June 2003 **Proceedings of the 2003 ACM SIGMOD internati nal conference n Management of data**

Full text available: pdf(292.76 KB) Additional Information: full citation, abstract, refer nces, index terms

A query optimizer compares alternative plans in its search space to find the best plan for a given query. Depending on the search space and the enumeration algorithm, optimizers vary in their compilation time and the quality of the execution plan they can generate. This paper describes a compilation time estimator that provides a quantified estimate of the optimizer compilation time for a given query. Such an estimator is useful for automatically choosing the right level of optimization in comme ...

77 XML query processing I: Composing XSL transformations with XML publishing views
Chengkai Li, Philip Bohannon, P. P. S. Narayan
June 2003 **Proceedings of the 2003 ACM SIGMOD international conference on Management of data**
Full text available: pdf(225.65 KB) Additional Information: full citation, abstract, references, citings, index terms

While the XML Stylesheet Language for Transformations (XSLT) was not designed as a query language, it is well-suited for many query-like operations on XML documents including selecting and restructuring data. Further, it actively fulfills the role of an XML query language in modern applications and is widely supported by application platform software. However, the use of database techniques to optimize and execute XSLT has only recently received atten ...

78 Query processing and optimization in Oracle Rdb
Gennady Antoshenkov, Mohamed Ziauddin
December 1996 **The VLDB Journal — The International Journal on Very Large Data Bases**, Volume 5 Issue 4
Full text available: pdf(92.62 KB) Additional Information: full citation, abstract, index terms

This paper contains an overview of the technology used in the query processing and optimization component of Oracle Rdb, a relational database management system originally developed by Digital Equipment Corporation and now under development by Oracle Corporation. Oracle Rdb is a production system that supports the most demanding database applications, runs on multiple platforms and in a variety of environments.

Keywords: Dynamic optimization, Optimizer, Query transformation, Relational database, Sampling

79 A benchmark of NonStop SQL release 2 demonstrating near-linear speedup and scaleup on large databases
Susanne Englert, Jim Gray, Terrye Kocher, Praful Shah
April 1990 **ACM SIGMETRICS Performance Evaluation Review , Proceedings of the 1990 ACM SIGMETRICS conference on Measurement and modeling of computer systems**, Volume 18 Issue 1
Full text available: pdf(251.80 KB) Additional Information: full citation, references, citings, index terms

80 Query processing techniques in the summary-table-by-example database query language
Gultekin Özsoyoğlu, Victor Matos, Meral Özsoyoğlu
December 1989 **ACM Transactions on Database Systems (TODS)**, Volume 14 Issue 4
Full text available: pdf(3.52 MB) Additional Information: full citation, abstract, r ferences, citings, index terms, review

Summary-Table-by-Example (STBE) is a graphical language suitable for statistical database applications. STBE queries have a hierarchical subquery structure and manipulate summary tables and relations with set-valued attributes. The hierarchical arrangement of STBE

queries naturally implies a tuple-by-tuple subquery evaluation strategy (similar to the nested loops join implementation technique) which may not be the best query processing strategy. In this paper we discuss the query ...

Results 61 - 80 of 200 Result page: previous 1 2 3 4 5 6 7 8 9 10 next

Terms used **sql outer join with**

Found **295** of **24,012** searched out of **86,143.**

Sort results by: relevance

Display results: expanded form

Save results to a Binder

Search Tips

☐ Open results in a new window

Try an Advanced Search

Try this search in The ACM Guide

Results 1 - 20 of 200 Result page: **1** 2 3 4 5 6 7 8 9 10 next

Best 200 shown

Relevance scale ☐▫▪■■

1 No regression algorithm for the enumeration of projections in SQL queries with joins and outer joins

Gautam Bhargava, Piyush Goel, Balakrishna R. Iyer

November 1995 **Proceedings of the 1995 conference of the Centre for Advanced Studies on Collaborative research**

Full text available: pdf(314.80 KB) Additional Information: full citation, abstract, references, citings, index terms

The execution time of an SQL query can be reduced significantly by considering different schedules for the operations specified in the query. The benefits of considering different schedules, a hallmark of strength in query optimization, are not usually exploited if a query contains projections along with binary operations. This paper presents a set of "no regression" algorithms that are capable of generating different schedules for the queries containing projections and binary operations.

2 Hypergraph based reorderings of outer join queries with complex predicates

Gautam Bhargava, Piyush Goel, Bala Iyer

May 1995 **ACM SIGMOD Record , Proceedings of the 1995 ACM SIGMOD international conference on Management of data**, Volume 24 Issue 2

Full text available: pdf(1.06 MB) Additional Information: full citation, abstract, references, citings, index terms

Complex queries containing outer joins are, for the most part, executed by commercial DBMS products in an "as written" manner. Only a very few reorderings of the operations are considered and the benefits of considering comprehensive reordering schemes are not exploited. This is largely due to the fact there are no readily usable results for reordering such operations for relations with *duplicates* and/or outer join predicates that are other than "*simple.*" Most previous approaches h ...

3 SQL query optimization: reordering for a general class of queries

Piyush Goel, Bala Iyer

June 1996 **ACM SIGMOD Record , Proceedings of the 1996 ACM SIGMOD international c nference n Management of data**, Volume 25 Issue 2

Full text available:

are no known methods to generate all equivalent reorderings for a SQL query containing joins, outer joins, and groupby aggregations. Consequently, some of the reorderings with significantly lower cost may be missed. Using hypergraph model and a set of novel identities, we propose a method to reorder a SQL query conta ...

4 Simplification of outer joins

Gautam Bhargava, Piyush Goel, Balakrishna R. Iyer

November 1995 **Proceedings of the 1995 conference of the Centre for Advanced Studies on Collaborative research**

Full text available: pdf(309.09 KB) Additional Information: full citation, abstract, references, citings, index terms

The removal of redundant outer joins is essential for the reassociation of outer joins with other binary operations. In this paper, we present a set of comprehensive algorithms that employ the properties of strong predicates along with the properties of SQL's projection, intersection, union and except operations in order to remove redundant outer joins from a complex query. For the purpose of query simplification, we generate additional projections by determining the keys. Our algorithm for gene ...

5 SilkRoute: A framework for publishing relational data in XML

Mary Fernández, Yana Kadiyska, Dan Suciu, Atsuyuki Morishima, Wang-Chiew Tan

December 2002 **ACM Transactions on Database Systems (TODS)**, Volume 27 Issue 4

Full text available: pdf(687.91 KB) Additional Information: full citation, abstract, references, citings, index terms

XML is the "lingua franca" for data exchange between interenterprise applications. In this work, we describe SilkRoute, a framework for publishing relational data in XML. In SilkRoute, relational data is published in three steps: the relational tables are presented to the database administrator in a canonical XML view; the database administrator defines in the XQuery query language a public, virtual XML view over the canonical XML view; and an application formulates an XQuery query over the publ ...

Keywords: XML, XML storage systems, XQuery

6 SchemaSQL: An extension to SQL for multidatabase interoperability

Laks V. S. Lakshmanan, Fereidoon Sadri, Subbu N. Subramanian

December 2001 **ACM Transactions on Database Systems (TODS)**, Volume 26 Issue 4

Full text available: pdf(435.89 KB) Additional Information: full citation, abstract, references, index terms, review

We provide a principled extension of SQL, called *SchemaSQL*, that offers the capability of uniform manipulation of data and schema in relational multidatabase systems. We develop a precise syntax and semantics of *SchemaSQL* in a manner that extends traditional SQL syntax and semantics, and demonstrate the following. (1) *SchemaSQL* retains the flavor of SQL while supporting querying of both data and schema. (2) It can be used to transform data in a database in a structure substa ...

Keywords: Information integration, SchemaSQL, multidatabase systems, restructuring views, schematic heterogeneity

7 Industrial sessions: database internals - II: Data densification in a relational database system

Abhinav Gupta, Sankar Subramanian, Srikanth Bellamkonda, Tolga Bozkaya, Nathan Folkert, Lei Sheng, Andrew Witkowski

June 2004 **Pr ceedings f the 2004 ACM SIGMOD internati nal conference on Management of data**

Full text available: pdf(227.07 KB) Additional Information: full citation, abstract, references

Data in a relational data warehouse is usually sparse. That is, if no value exists for a given combination of dimension values, no row exists in the fact table. Densities of 0.1-2% are very common. However, users may want to view the data in a dense form, with rows for all combination of dimension values displayed even when no fact data exists for them. For example, if a product did not sell during a particular time period, users may still want to see the product for that time period with zero s ...

8 ANSI SQL hierarchical processing can fully integrate native XML

Michael M. David

March 2003 **ACM SIGMOD Record**, Volume 32 Issue 1

Full text available: pdf(162.35 KB) Additional Information: full citation, abstract, references

Most SQL-based XML vendor support is through interoperation and not integration. One reason for this is that XML is inherently hierarchical and SQL is supposedly not. This paper demonstrates how ANSI SQL along with its relational Cartesian product model can naturally perform complete and flexible hierarchical query processing. With this ANSI SQL inherent hierarchical processing capability, native XML data can be fully and seamlessly integrated into SQL processing and operated on at a full hierar ...

9 Research sessions: spatial data: Joining interval data in relational databases

Jost Enderle, Matthias Hampel, Thomas Seidl

June 2004 **Proceedings of the 2004 ACM SIGMOD international conference on Management of data**

Full text available: pdf(552.80 KB) Additional Information: full citation, abstract, references

The increasing use of temporal and spatial data in present-day relational systems necessitates an efficient support of joins on interval-valued attributes. Standard join algorithms do not support those data types adequately, whereas special approaches for interval joins usually require an augmentation of the internal access methods which is not supported by existing relational systems. To overcome these problems we introduce new join algorithms for interval data. Based on the Relational Interval ...

10 Efficient evaluation of XML middle-ware queries

Mary Fernandez, Atsuyuki Morishima, Dan Suciu

May 2001 **ACM SIGMOD Record , Proceedings of the 2001 ACM SIGMOD international conference on Management of data**, Volume 30 Issue 2

Full text available: pdf(414.15 KB) Additional Information: full citation, abstract, references, citings, index terms

We address the problem of efficiently constructing materialized XML views of relational databases. In our setting, the XML view is specified by a query in the declarative query language of a middle-ware system, called SilkRoute. The middle-ware system evaluates a query by sending one or more SQL queries to the target relational database, integrating the resulting tuple streams, and adding the XML tags. We focus on how to best choose the SQL queries, without having control over the target RDBM ...

11 On random sampling over joins

Surajit Chaudhuri, Rajeev Motwani, Vivek Narasayya

June 1999 **ACM SIGMOD Rec rd , Proceedings f the 1999 ACM SIGMOD internati nal c nference n Management of data**, Volume 28 Issue 2

Full text available: pdf(1.65 MB) Additional Information: full citation, abstract, refer nces, citings, ind x terms

A major bottleneck in implementing sampling as a primitive relational operation is the inefficiency of sampling the output of a query. It is not even known whether it is possible to generate a sample of a join tree without first evaluating the join tree completely. We undertake a detailed study of this problem and attempt to analyze it in a variety of settings. We present theoretical results explaining the difficulty of this problem and setting limits on the efficiency that can be achieved. ...

12 Full papers: Horizontal aggregations for building tabular data sets

Carlos Ordonez

June 2004 **Proceedings of the 9th ACM SIGMOD workshop on Research issues in data mining and knowledge discovery**

Full text available: pdf(183.18 KB) Additional Information: full citation, abstract, references

In a data mining project, a significant portion of time is devoted to building a data set suitable for analysis. In a relational database environment, building such data set usually requires joining tables and aggregating columns with SQL queries. Existing SQL aggregations are limited since they return a single number per aggregated group, producing one row for each computed number. These aggregations help, but a significant effort is still required to build data sets suitable for data mining pu ...

13 On optimizing an SQL-like nested query

Won Kim

September 1982 **ACM Transactions on Database Systems (TODS)**, Volume 7 Issue 3

Full text available: pdf(1.79 MB) Additional Information: full citation, abstract, references, citings, index terms

SQL is a high-level nonprocedural data language which has received wide recognition in relational databases. One of the most interesting features of SQL is the nesting of query blocks to an arbitrary depth. An SQL-like query nested to an arbitrary depth is shown to be composed of five basic types of nesting. Four of them have not been well understood and more work needs to be done to improve their execution efficiency. Algorithms are developed that transform queries involving these basic ty ...

Keywords: aggregate function, divide, join, nested query, predicate, relational database

14 Translation with optimization from relational calculus to relational algebra having aggregate functions

Ryohei Nakano

December 1990 **ACM Transactions on Database Systems (TODS)**, Volume 15 Issue 4

Full text available: pdf(2.77 MB) Additional Information: full citation, abstract, references, citings, index terms

Most of the previous translations of relational calculus to relational algebra aimed at proving that the two languages have the equivalent expressive power, thereby generating very complicated relational algebra expressions, especially when aggregate functions are introduced. This paper presents a rule-based translation method from relational calculus expressions having both aggregate functions and null values to optimized relational algebra expressions. Thus, logical optimization is carrie ...

15 A history and evaluation of System R

Donald D. Chamberlin, Morton M. Astrahan, Michael W. Blasgen, James N. Gray, W. Frank King, Bruce G. Lindsay, Raymond Lorie, James W. Mehl, Thomas G. Price, Franco Putzolu, Patricia Griffiths Selinger, Mario Schkolnick, Donald R. Slutz, Irving L. Traiger, Bradford W. Wade, Robert A. Yost

October 1981 **C mmunicati ns f the ACM**, Volume 24 Issue 10

Full text available: pdf(1.55 MB) Additional Information: full citation, abstract, references, citings, index terms

System R, an experimental database system, was constructed to demonstrate that the usability advantages of the relational data model can be realized in a system with the complete function and high performance required for everyday production use. This paper describes the three principal phases of the System R project and discusses some of the lessons learned from System R about the design of relational systems and database systems in general.

Keywords: access path selection, authorization, compilation, database management systems, locking, recovery, relational model

16 Implementation of magic-sets in a relational database system

Inderpal Singh Mumick, Hamid Pirahesh

May 1994 **ACM SIGMOD Record , Proceedings of the 1994 ACM SIGMOD international conference on Management of data**, Volume 23 Issue 2

Full text available: pdf(1.34 MB) Additional Information: full citation, abstract, references, citings, index terms

We describe the implementation of the magic-sets transformation in the Starburst extensible relational database system. To our knowledge this is the first implementation of the magic-sets transformation in a relational database system. The Starburst implementation has many novel features that make our implementation especially interesting to database practitioners (in addition to database researchers). (1) We use a cost-based heuristic for determining join orders (sips) bef ...

17 Efficiently publishing relational data as XML documents

Jayavel Shanmugasundaram, Eugene Shekita, Rimon Barr, Michael Carey, Bruce Lindsay, Hamid Pirahesh, Berthold Reinwald

September 2001 **The VLDB Journal — The International Journal on Very Large Data Bases**, Volume 10 Issue 2-3

Full text available: pdf(216.67 KB) Additional Information: full citation, abstract, citings, index terms

XML is rapidly emerging as a standard for exchanging business data on the World Wide Web. For the foreseeable future, however, most business data will continue to be stored in relational database systems. Consequently, if XML is to fulfill its potential, some mechanism is needed to publish relational data as XML documents. Towards that goal, one of the major challenges is finding a way to efficiently structure and tag data from one or more tables as a hierarchical XML document. Different alterna ...

Keywords: Publishing, Relational databases, XML

18 A critique of the SQL database language

C. J. Date

November 1984 **ACM SIGMOD Record**, Volume 14 Issue 3

Full text available: pdf(2.38 MB) Additional Information: full citation, abstract, references

The ANS Database Committee (X3H2) is currently at work on a proposed standard relational database language (RDL), and has adopted as a basis for that activity a definition of the "structured query language" SQL from IBM [10]. Moreover, numerous hardware and software vendors (in addition to IBM) have already released or at least announced products that are based to a greater or lesser extent on the SQL language as defined by IBM. There can thus be little doubt that the importance of that language ...

19 MySQL Introduction
David Axmark, Michael Widenius
November 1999 **Linux J urnal**

Full text available: html(35.14 KB) Additional Information: full citation, abstract, index terms

A look at the MySQL database--where it's been, where it is now, and where it's going

20 Advanced XML technologies and applications: Honey, I shrunk the XQuery!: an XML algebra optimization approach
Xin Zhang, Bradford Pielech, Elke A. Rundesnteiner
November 2002 **Proceedings of the 4th international workshop on Web information and data management**

Full text available: pdf(375.44 KB) Additional Information: full citation, abstract, references, citings, index terms

A lot of work is being done in the database community on mapping of XML data into and out of relational database systems, specifically, the query processing over such data using XQuery. We discuss our solution, the XML Algebra Tree (XAT), as part of our larger XML management system called Rainbow.Rainbow uses XQuery to describe the loading and extracting of XML data into relational systems and also for the execution of queries against pre-defined XML views of that stored data. The XML algebra tr ...

Keywords: XML, XQuery, algebra, operator, optimization, relational

Results 1 - 20 of 200 Result page: **1** 2 3 4 5 6 7 8 9 10 next

Terms used
sql outer join with test or **diagnose database queries**

Found **3** of **295** searched out of **295**.

Sort results by: relevance
Display results: expanded form
Save results to a Binder
Search Tips
☐ Open results in a new window

Try an Advanced Search
Try this search in The ACM Guide

Results 1 - 3 of 3

Relevance scale ☐■■■■

1 Research sessions: consistency and availability: Relaxed currency and consistency: how to say "good enough" in SQL

Hongfei Guo, Per-Åke Larson, Raghu Ramakrishnan, Jonathan Goldstein

June 2004 **Proceedings of the 2004 ACM SIGMOD international conference on Management of data**

Full text available: pdf(606.81 KB) Additional Information: full citation, abstract, references, citings

Despite the widespread and growing use of asynchronous copies to improve scalability, performance and availability, this practice still lacks a firm semantic foundation. Applications are written with some understanding of which queries can use data that is not entirely current and which copies are "good enough"; however, there are neither explicit requirements nor guarantees. We propose to make this knowledge available to the DBMS through explicit currency and consistency (C&C) constraints in qu ...

2 Extensible/rule based query rewrite optimization in Starburst

Hamid Pirahesh, Joseph M. Hellerstein, Waqar Hasan

June 1992 **ACM SIGMOD Record , Proceedings of the 1992 ACM SIGMOD international conference on Management of data**, Volume 21 Issue 2

Full text available: pdf(1.30 MB) Additional Information: full citation, abstract, references, citings, index terms

This paper describes the Query Rewrite facility of the Starburst extensible database system, a novel phase of query optimization. We present a suite of rewrite rules used in Starburst to transform queries into equivalent queries for faster execution, and also describe the production rule engine which is used by Starburst to choose and execute these rules. Examples are provided demonstrating that these Query Rewrite transformations lead to query execution time improvements of orders of magni ...

3 Multi Relational Data Mining (MRDM): Scalability and efficiency in multi-relational data mining

Hendrik Blockeel, Michèle Sebag

July 2003 **ACM SIGKDD Explorations Newsletter**, Volume 5 Issue 1

Full text available: pdf(1.61 MB) Additional Information: full citation, abstract, r ferences, citings

Efficiency and Scalability have always been important concerns in the field of data mining, and are even more so in the multi-relational context, which is inherently more complex. The issue has been receiving an increasing amount of attention during the last few years, and quite a number of theoretical results, algorithms and implementations have been presented

that explicitly aim at improving the efficiency and Scalability of multi-relational data mining approaches. With this article we attempt ...

Results 1 - 3 of 3